Grid Computing Technology, the OAIS Reference Model, and Persistent Archive Environments

Bruce R. Barkstrom and
David E. Cordner

Atmospheric Sciences Data Center
NASA Langley Research Center

Outline

- Challenges with Current Data
  - Requirements for Expert Knowledge
  - Data Management
- The Role of Commodity Computing and Grid Technology
- Help from the OAIS Reference Model
- Preservation Challenges
  - Hardware Perishes – Data Needs Immortality
  - Human Knowledge Requires Human Communities
  - Overcoming Death and Taxes

Challenges with Current Data

- Conventional View of Challenges
  - Large Volumes: ~10 PB in current DAACs
  - Complex Formats:
    - But data are still “images”
    - HDF manages – but isn’t universally accepted by user community
  - Production: Delimited by Levels –
    0 -> 1, 1 -> 2, 2 -> 3
  - Cost of Preservation: Attributed to missions
    - When mission funding disappears, so does preservation

Requirements for Expert Knowledge

- Measurements Come From Complex Physical Chains
  - Instruments are complex
    - “Calibration” should be inverse of measurement
  - Satellite sampling is intricate
    - Instrument sampling compounds orbit sampling
  - Reduction to geophysical parameters requires rigorous derivation
- Stored Data is Repository of Expert Human Knowledge

Data Management – I. Production

- Data Production can be complex
  - Production topology may not be simple

    0 -> 1, 1 -> 2, 2 -> 3
  - Production flow may be discrete and intermittent
  - Validation usually creates reentrant flows
  - ASDC has two production examples (MISR and CERES) each with more than 1M SLOC

Data Management – II. Users

- ECS Design predicated on small orders of discrete files to fairly large user community
  - Suitable for sample images, case studies
  - Requires caches for field experiment groupings – and needs to catch data on way from production to archive
- Other user communities need different kinds of access
  - Large scale climate work either requires validated L3 data (with complex rework production flow) or content-based data streaming
    - 105,000 files and 30 TB of CERES data for examining 12 years of L2 data
  - Large-scale, interdisciplinary climate work requires coordination of data flows between data centers
    - Investigation of storms between microwave and radiation may require long time series of physically synchronous intercomparisons
  - Time series investigations may require database subsets
- Most users are not well-prepared to handle multi-TB data sets

The Role of Commodity Computing and Grid Technology

- Data uses seem well-suited to “one-file per CPU” computation
  - Not many CPU’s per large array needed for models
- Commodity computing reduces HW costs
  - Clusters well suited to high-throughput data processing
- Grid computing can make it easier to balance data flows and coordinated computing between centers

Help From the OAIS Reference Model

- Open Archive Information Systems (OAIS) Reference Model
  - ISO standard providing description of archive functions and data flows
- Can help produce a “flow-based” architecture
  - Allows identification of automatable data management workflows
  - Good basis for standard protocols to help with modularity and survivable components

OAIS Reference Model Flows

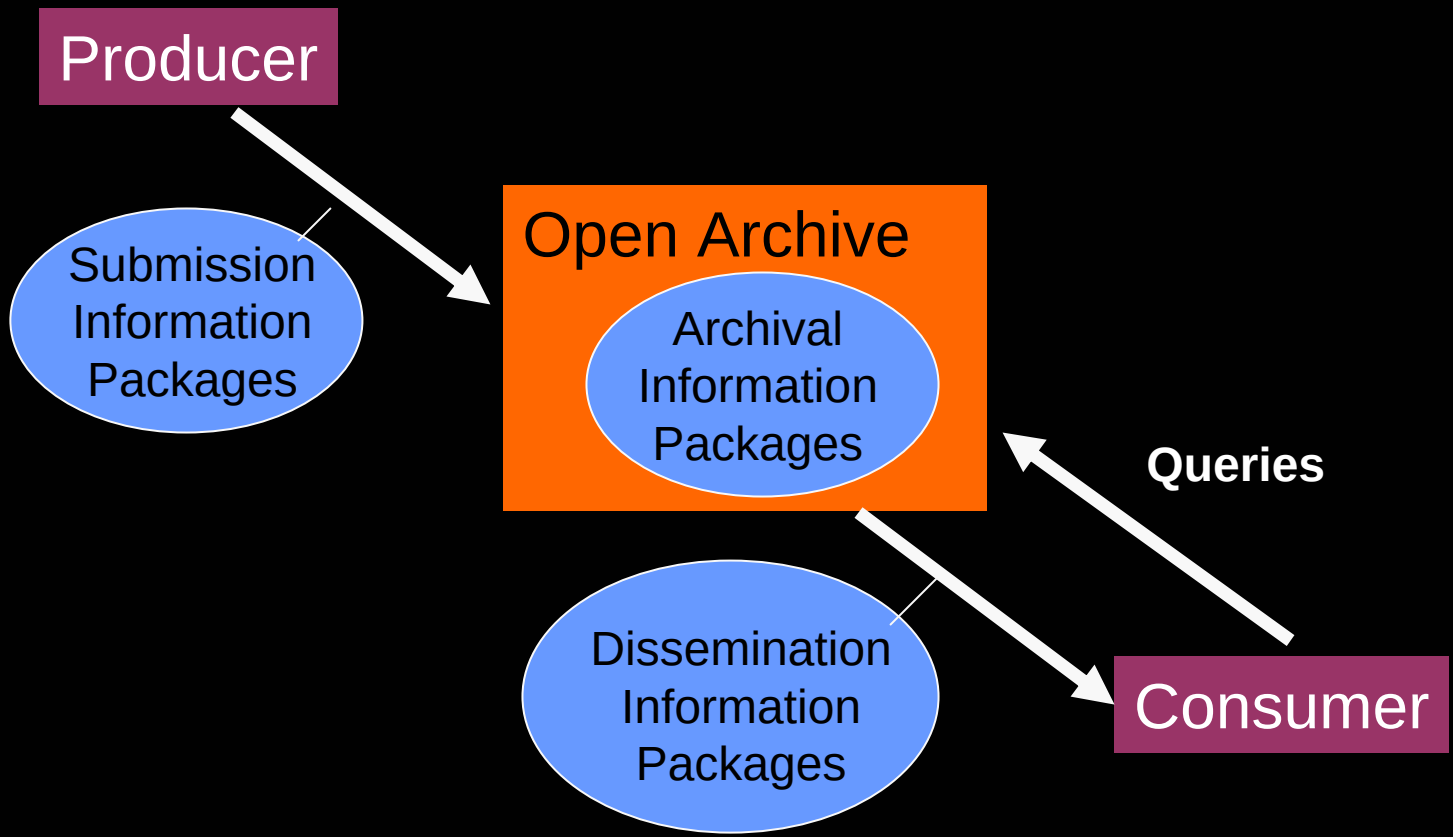

Preservation Challenges

- Basic Challenges of Preservation are “Sociological”
  - Knowledge is created by human communities, not by hardware or software
  - Social boundaries create real barriers to preserving created knowledge or to creating new knowledge
    - Tribal vocabularies and world views
    - Tribal customs and power relationships

Hardware Perishes – Data Needs Immortality

- Conventional view seems to assume preserving media preserves knowledge
- Actually, hardware is obsolete in 5 years
- Software creators and vendors are perishable organizations
- Major reason for migrating data is reducing cost by taking advantage of new hardware/software capability

Human Knowledge Requires Human Communities

- Archives and data centers need to assist in preserving community knowledge
  - Serious requirement to gather calibration and algorithm knowledge before producer teams disband
- Need to visualize knowledge communities as extending beyond mission and agency boundaries
  - Science teams are often academies of disciplinary knowledge that have much longer lives than particular missions
  - Science team work can be much more expensive if data access is restricted

Overcoming Death and Taxes

- Largest threats to knowledge loss are social
  - IT Security (threat to chain-of-custody)
  - Operator Error
  - Funding
- Future archives
  - Need to avoid errors
    - Data will die if error rate exceeds ~10^{-5} per year
  - Need to overcome institutional and disciplinary boundaries
    - Knowledge will die if resources not available, may want to consider ‘Open Source Archives’ and serious interagency cooperation

Hurricane Isabel: What We Knew When and What We Did – Friday, Sept. 12

- First Indicators of Isabel as Cat 5 Hurricane in Caribbean on Friday, Sept. 12
- ASDC Head requested emergency tape evacuation procedure from System Engineer – received late on Friday afternoon
- ASDC Head notified Atmospheric Sciences Competency Director Sunday evening, noting possibility of disaster evacuation – Director concurs

Hurricane Isabel: What We Knew When and What We Did – Friday, Sept. 12

- First Indicators of Isabel as Cat 5 Hurricane in Caribbean on Friday, Sept. 12
- ASDC Head requested emergency tape evacuation procedure from System Engineer – received late on Friday afternoon
- ASDC Head notified Atmospheric Sciences Competency Director Sunday evening, noting possibility of disaster evacuation – Director concurs

Monday, Sept. 15, 2003

- National Hurricane Center storm track and strength constant over last 36 hours – Cat 5 until landfall, with storm track overhead
- Landfall expected Thursday, Sept. 18 – need to evacuate tapes by Tuesday to get safely to Ashland, VA before evacuation traffic
- Staff meeting early morning – ASDC Head decides to order Iron Mountain trucks
- Trucks ordered about 1 pm – cost < $16k
- Production halted; systems start shut-down

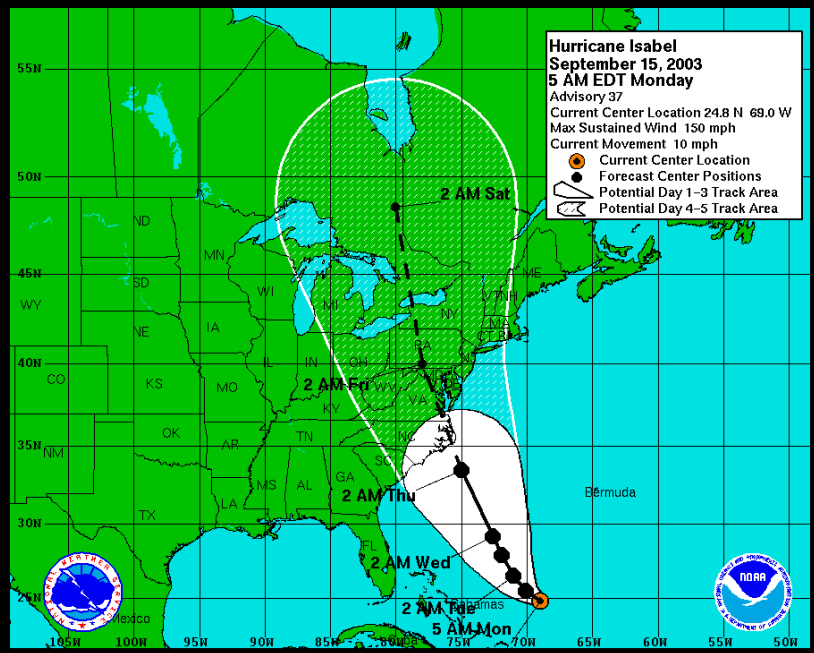

Tuesday, Sept. 16

- National Hurricane Center storm track now significantly west of LaRC, storm intensity downgraded to high Cat 3
- ASDC Head met with AtSC Director – danger sufficiently down to rescind order for trucks
- Trucks show up about 9:30 am – Iron Mountain staff given tour and posters (Decision irrevocable – if storm surge 25 ft, will lose tapes and other equipment)
- Production restarted

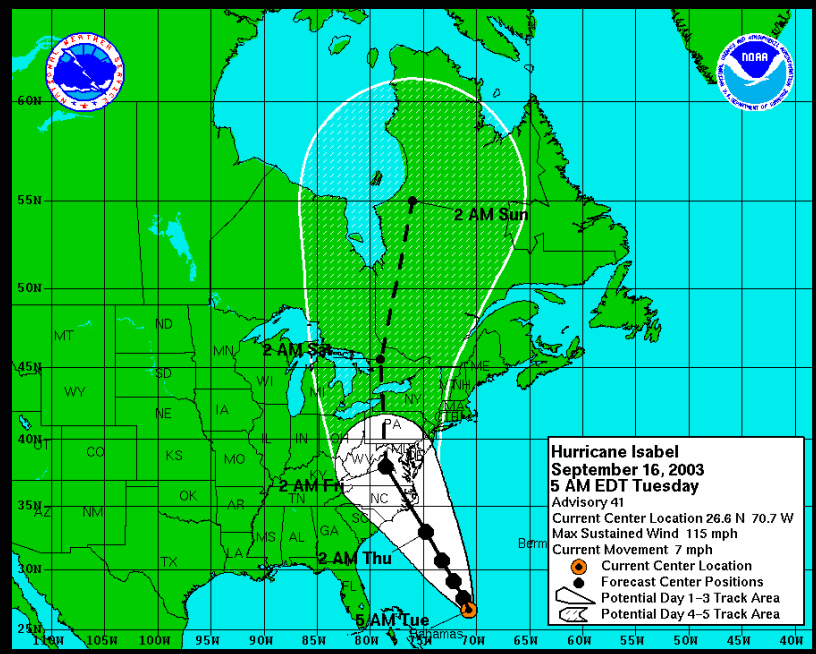

Thursday, Sept. 18

- Hurricane landfall mid-afternoon
- 6:15 am – first reasonable forecast of record storm surge for stations near mouth of Chesapeake Bay
- LaRC closed
- Power lost in Williamsburg about 2 pm – last power or reliable phone service for 7 days
- Storm closes in – wind and rain, with occasional torrential rain bursts and loud tree noises

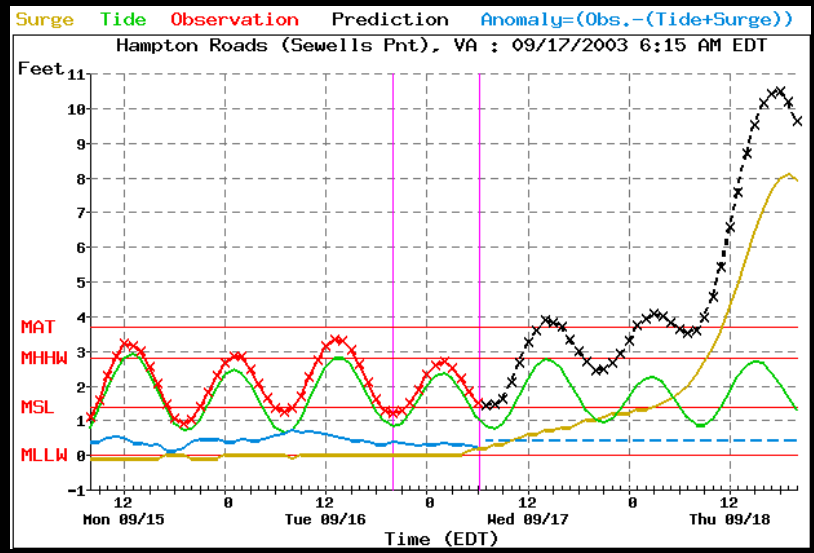

LaRC Storm Surge

- Isabel storm surge record high – higher than 1933 hurricane in Poquoson
- Isabel only Cat 2 at Langley – storm surge still 10 feet above MLLW
- Surge rise at rate of 1 inch per minute – cars float at 2 feet: mortal danger within twenty minutes of water starting to rise
- With Cat 5 storm, 20 to 25 foot surge possible – base of ASDC about 10 feet above MLLW

A Lost Weekend – Sept. 19-21

Williamsburg – 35 miles from LaRC: Microbursts topple trees onto houses; Trees down power lines; 1.8 Million residents of Hampton Roads without electrical power; Gas not available; Stoplights not operating.

Risk Analysis and Mitigation

- Standard Procedure for Insurance Valuation
- Steps:
  - Assess sources of value
  - Identify threats
  - Assess probability of threat and of loss
  - Mitigate risk through avoidance, mitigation, insurance

Probability of Loss

Threat	Loss Probability per Year
Hurricane – Cat II or greater	0.02
Hurricane – Cat V	0.005
Tornados, Aircraft, Earthquakes, Nuclear Reactors, Terrorists	0.005
IT Attacks	**0.1**

Probability of Survival

- Survival for 200 years (archival standard) is hard

$$P = (1 - \varepsilon)^N$$

  - P is probability of surviving N years
  - ε is probability of loss per year

  – If $\varepsilon \sim 0.1$ per year and $N \sim 200$, $P \sim 10^{-10}$
  Long Odds

- Lesson: Store data off-site, off-line

Derived Requirement

- Reduce Probability of Loss
- Corollaries:
  - Simplify systems to reduce errors
  - Diversify risk – avoid single failure points; Replicate data and system implementations
  - Reduce probability of operator error – Practice operations and installations (even during design)

Development Costs and Operations Costs

- Model – ASDC LaTIS Data System
  - 100,000 SLOC
  - ~1/2 PB of data
- Use commercial software cost est. tool
  - ~2 years, ~$10M for development
  - 5 years of maintenance and operations after delivery
- **Conclusion:**
  - **software maintenance and operations are 60 % of total cost**
  - **development only 40 % of cost**

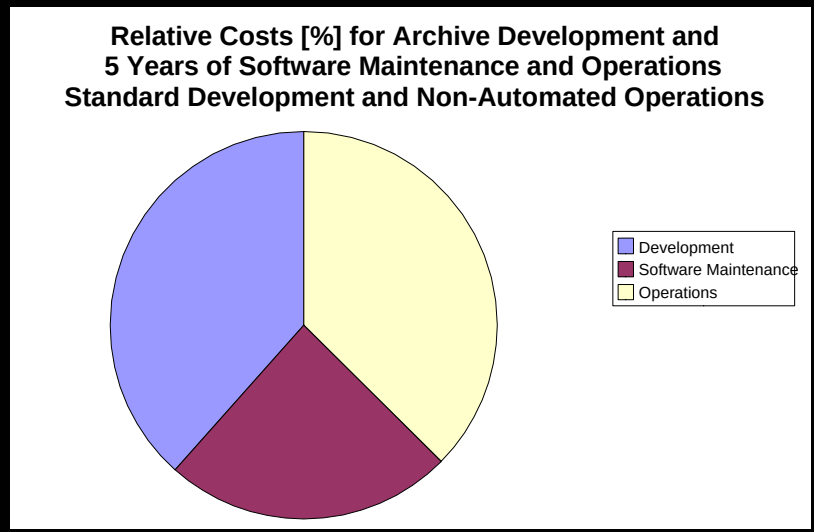

Derived Requirements

- Design for Automation and Low Defect Rate
- Corollaries:
  - Pay more attention to workflow than to functionality in architecture and design
  - Concentrate on measures that prevent errors
    REWORK IS EXPENSIVE
  - Use Open Source and Commodity Computing to reduce costs
  - Have developers practice installation and evolutionary upgrades to their systems

Users as Tribal Communities

- Users are members of “tribes”;
  So are managers
  - Distinct tribal vocabularies
  - Distinct tribal world views of data
  - Distinct tribal customs
- Tribes evolve
  - Vocabularies and concepts change
  - Managers subject to “management fashions”
    (for which there is a theory)

Some Signs of Hope

- Locally Autonomous Federations Work
  - Sharing resources primarily with trusted partners reduces probability of free loading
  - Potential for reducing managerial overhead
  - Need managerial wisdom in HQ organizations
- Reference Models Can Reduce Design Work and Produce Good Systems

Summary Recommendations

- Simplify
- Reduce Defects
- Design-in Automation
- Practice Operations
- Use Federated Systems – Not Imperial
- Embrace Change